

What is claimed is:

1. A file storage system comprising:

a gateway service adapted to service requests from clients for accessing files having

5 respective file identifiers;

a storage service adapted to provide permanent storage of file data corresponding to the files; and

a metadata service adapted to maintain metadata corresponding to the files, a portion of the metadata correlating respective locations of the files in the storage service with the respective
10 file identifiers,

the gateway service communicating with the metadata service to determine a location of the bitfile data in the storage service in accordance with the file identifier, and with the storage service for providing clients access to the bitfile data in accordance with the determined location, whereby the clients need only provide the file identifier to access the file from the determined
15 location.

2. A file storage system according to claim 1, wherein the storage service comprises a plurality of storage devices, the location of the bitfile data being associated with the plurality of storage devices.

3. A file storage system according to claim 2, wherein the storage service is adapted so that the number of storage devices is scalable.

4. A file storage system according to claim 2, wherein the storage service and the metadata service are adapted so that the location of the bitfile data corresponding to a certain one of the files is associated with more than one of the plurality of storage devices.

5

5. A file storage system according to claim 4, wherein the more than one of the plurality of storage devices are determined in accordance with a redundancy scheme.

6. A file storage system according to claim 5, wherein the gateway service is adapted to provide access to an operable device of the more than one of the plurality of storage devices rather than a failed device of the more than one of the plurality of storage devices in accordance with a client request associated with the certain one of the files and with the location determined from the metadata service.

7. A file storage system according to claim 4, wherein the more than one plurality of storage devices are determined in accordance with a striping scheme.

8. A file storage system according to claim 1, wherein the gateway service comprises a plurality of gateway service nodes, the gateway service nodes being adapted to uniformly communicate with the metadata service, the storage service and the clients.

9. A file storage system according to claim 8, further comprising a load balancer adapted to distribute the requests from the clients to the plurality of gateway service nodes.

10. A file storage system according to claim 8, wherein the gateway service nodes provide substantially the same performance in servicing the client file requests.

11. A file storage system according to claim 8, wherein the gateway service is adapted so that the number of gateway service nodes is scalable.

12. A file storage system according to claim 1, wherein the metadata service comprises a plurality of metadata nodes, the metadata nodes being adapted to uniformly communicate with the gateway service and to maintain the metadata.

13. A file storage system according to claim 1, wherein the metadata service is adapted to implement a hierarchical storage management scheme.

14. A file storage system according to claim 2, wherein the metadata service is adapted to implement a hierarchical storage management scheme, a first set of the plurality of storage devices having a first capacity and performance characteristic, a second set of the plurality of storage devices having a second capacity and performance characteristic different than the first characteristic, the metadata service cooperating with the storage service to manage a migration of bitfile data from the first set to the second set of storage devices, whereby clients requesting

access to moved files corresponding to the migrated bitfile data need only supply the file identifiers corresponding to the moved files.

15. A file storage system according to claim 1, wherein the clients comprise one of at least one
5 of a NFS client, a CIFS client, a HTTP client, and a FTP client, the gateway service being further adapted to service the client requests in accordance with the respective protocols used by the clients.

16. A file storage system comprising:

10 a scalable number of gateway servers each adapted to service requests from clients for accessing files having respective file identifiers;

a scalable number of storage servers each adapted to provide permanent storage of bitfile data corresponding to the files; and

15 a metadata service adapted to maintain metadata corresponding to the files, a portion of the metadata correlating respective locations of the files in the storage service with the respective file identifiers,

the gateway servers each communicating with the metadata service to determine a location of the bitfile data in the storage servers in accordance with the file identifier, and with the bitfile storage servers for providing clients access to the bitfile data in accordance with the
20 determined location, whereby a client who accessed the file when the scalable number comprised a first value need only provide the same file identifier to access the file when the scalable number comprises a second value different than the first value.

17. A method for providing shared access to stored files, comprising:

servicing requests from clients for accessing files having respective file identifiers;

providing permanent storage of file data corresponding to the files in a first storage

5 device; and

maintaining metadata corresponding to the files in a second storage device different than the first device, a portion of the metadata correlating respective locations of the files in the first storage device with the respective file identifiers,

determining a location of the bitfile data in the first device in accordance with the file
10 identifier from the metadata maintained in the second device; and

providing clients access to the bitfile data from the first device in accordance with the determined location, whereby the clients need only provide the file identifier to access the file from the determined location.

15 18. A method according to claim 17, wherein the first storage device comprises a number of storage devices, the location of the bitfile data being associated with the number of storage devices.

19. A method according to claim 18, wherein the maintaining step includes accounting for a
20 scaling of the number of storage devices.

20. A method according to claim 18, wherein the maintaining step includes correlating the location of the bitfile data corresponding to a certain one of the files being associated with more than one of the number of storage devices.

5 21. A method according to claim 20, further comprising determining the more than one of the number of storage devices in accordance with a redundancy scheme.

22. A method according to claim 21, further comprising providing access to an operable device of the more than one of the number of storage devices rather than a failed device of the more than one of the number of storage devices in accordance with a client request associated with the certain one of the files and with the determined location.

23. A method according to claim 20, further comprising determining the more than one of the number of storage devices in accordance with a striping scheme.

24. A method according to claim 17, wherein the servicing step includes servicing the requests in a uniform manner by a plurality of gateway service nodes.

25. A method according to claim 24, further comprising distributing the requests from the clients to the plurality of gateway service nodes.

26. A method according to claim 24, wherein the gateway service nodes provide substantially the same performance in servicing the client file requests.

27. A method according to claim 17, wherein the servicing step includes servicing the requests
5 in accordance with at least one of a NFS protocol, a CIFS protocol, a HTTP protocol, and a FTP protocol.

28. A method according to claim 17, further comprising:

migrating certain of the bitfile data from the first device to a third storage device different
10 than the first and second storage devices; and

updating the metadata corresponding to the migrated bitfile data to reflect the migrated
location of the migrated bitfile data,

whereby the clients need only provide the file identifier associated with the migrated
bitfile data to access the file from the migrated location

29. A method according to claim 28, further comprising:

removing the bitfile data from old locations in the first storage device.